

EL764085728

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION FOR LETTERS PATENT

**On-Disk File Format for a Serverless Distributed File
System**

Inventor(s):
William J. Bolosky
Gerald Cermak
Atul Adya
John R. Douceur

ATTORNEY'S DOCKET NO. MS1-733US

TECHNICAL FIELD

This invention relates to serverless distributed file systems, and particularly to formats of files stored in serverless distributed file systems.

BACKGROUND

File systems manage files and other data objects stored on computer systems. File systems were originally built into a computer's operating system to facilitate access to files stored locally on resident storage media. As computers became networked, some file storage capabilities were offloaded from individual user machines to special storage servers that stored large numbers of files on behalf of the user machines. When a file was needed, the user machine simply requested the file from the server. In this server-based architecture, the file system was extended to facilitate management of and access to files stored remotely at the storage server over a network.

One problem that arises in distributed file systems concerns storage of identical files on the server. While some file duplication normally occurs on an individual user's personal computer, duplication unfortunately tends to be quite prevalent on networks where a server centrally stores the contents of multiple personal computers. For example, with a remote boot facility on a computer network, each user boots from that user's private directory on a file server. Each private directory thus ordinarily includes a number of files that are identical to files on other users' directories. Storing the private directories on traditional file systems consumes a great amount of disk and server file buffer cache space. From a storage management perspective, it is desirable to minimize file duplication to reduce the amount of wasted storage space used to store redundant files.

1 However, any such efforts need to be reconciled with the file system that tracks
2 the multiple duplicated files on behalf of the associated users.

3 To address the problems associated with storing multiple identical files on a
4 computer, Microsoft developed a single instance store (SIS) system that is
5 packaged as part of the Windows 2000 operating system. The SIS system reduces
6 file duplication by automatically identifying common identical files of a file
7 system, and then merging the files into a single instance of the data. One or more
8 logically separate links are then attached to the single instance to represent the
9 original files to the user machines. In this way, the storage impact of duplicate
10 files on a computer system is greatly reduced.

11 Today, file storage is migrating toward a model in which files are stored on
12 various networked computers, rather than on central storage server. The serverless
13 architecture poses new challenges to file systems. One particular challenge
14 concerns managing files that are distributed over many different computers in a
15 manner that allows a user to quickly access a file, verify that it is indeed the
16 requested file, and read/write that file, all while insuring that the files are stored
17 and accessed in a secure way that prevents access by non-authorized users.

18 The invention addresses these challenges and provides solutions that are
19 effective for distributed file systems, and in particular, serverless distributed file
20 systems.

21 22 SUMMARY

23 A file format for a serverless distributed file system is composed of two
24 parts: a primary data stream and a metadata stream. The data stream contains a
25 file that is divided into multiple blocks. Each block is encrypted using a

1 symmetric cipher (e.g., RC4) and a hash of the block as the encryption key. The
2 metadata stream contains a header, a structure for indexing the encrypted blocks in
3 the primary data stream, and some user information.

4 The indexing tree structure defines leaf nodes for each of the blocks. Each
5 leaf node consists of an access value used for decryption of the associated block
6 and a verification value used to independently verify the encrypted block
7 independently of other blocks. In one implementation, the access value is formed
8 by hashing the file block and encrypting the resultant hash value using a
9 symmetric cipher and a randomly generated key. The key is then encrypted using
10 an asymmetric cipher (e.g., RSA) and the user's public key as the encryption key.
11 The verification value is formed by hashing the associated encrypted block using a
12 one-way hash function (e.g., SHA).

13 Depending on the size of the file, the indexing structure may include
14 intermediate nodes formed by grouping the leaf nodes into tree blocks and
15 computing hash values of each tree block. These intermediate nodes can again be
16 segmented into blocks and each block hashed to form the next nodes. This can be
17 repeated as many times as desired until reaching a root node. The root node is
18 then hashed, and the hash value is used along with the metadata header and user
19 information to produce a verification value for the entire file. In one
20 implementation, the whole-file verification value is signed with a user's signature.
21 Alternatively, a file may be constructed without such signatures.

22 The file format supports verification of individual file blocks without
23 knowledge of the randomly generated key or any user keys. To verify a block of
24 the file, the file system optionally evaluates the signature on whole file verification
25 value (if one exists), checks that the whole-file verification value matches the hash

1 of the root block, metadata header and user information and then traverses the tree
2 to the appropriate leaf node associated with a target block to be verified. The file
3 system hashes the target block and if the hash matches the access value contained
4 in the leaf node, the block is authentic.

5 The file format further supports reading from and writing to individual
6 blocks without interfering with other blocks. The file format is also conducive for
7 sparse files that have vast areas of non-data.

8 9 **BRIEF DESCRIPTION OF THE DRAWINGS**

10 The same numbers are used throughout the document to reference like
11 components and/or features.

12 Fig. 1 is an illustration of a networked computing system that implemented
13 a serverless distributed file system.

14 Fig. 2 is a block diagram of logical components implemented at each of the
15 computers in the computing system of Fig. 1.

16 Fig. 3 is a block diagram of a computer that may be used to implement a
17 computer in the computing system of Fig. 1.

18 Fig. 4 illustrates a file format for a file stored in the serverless distributed
19 file system. The file format includes a data stream and a metadata stream.

20 Fig. 5 shows a file that is segmented and encrypted to form the data stream
21 of the file format.

22 Fig. 6 shows an indexing structure to index the file segments. The indexing
23 structure forms part of the metadata stream.

24 Fig. 7 illustrates a user key list that identifies users with privileges to access
25 the file.

1 Fig. 8 is a flow diagram of a process for constructing a file according to the
2 file format of Fig. 4.

3 Fig. 9 is a flow diagram of a process for verifying individual blocks of the
4 file.

5 Fig. 10 is a flow diagram of a process for reading individual blocks of the
6 file.

7 Fig. 11 is a flow diagram of a process for writing to a file block.

8 Fig. 12 is a flow diagram of a process for producing a signed manifest of
9 changes made to one or more files.

10 Fig. 13 illustrates a signed manifest and exemplary contents therein.

11 DETAILED DESCRIPTION

12 The following discussion is directed to a file format used in a distributed
13 file system, and to techniques for managing access to and verification of files
14 using the file format. The file system is described in the context of a symbiotic,
15 serverless, distributed file system that runs on multiple networked computers and
16 stores files across the computers rather than on a central server or cluster of
17 servers. The symbiotic nature implies that the machines cooperate but do not
18 completely trust one another.
19

20 The file system does not manage the storage disk directly, but rather relies
21 on existing file systems on local machines, such as those file systems integrated
22 into operating systems (e.g., the Windows NT® file system). The file system
23 described herein assumes a level at which multi-stream files can be stored as the
24 underlying local storage model.
25

1 While the file system is described in the context of storing "files", it should
2 be noted that other types of storable data can be stored in the file system. The
3 term "file" is used for discussion purposes and is intended to include data objects
4 or essentially any other storage subject matter that may not be commonly
5 characterized as a "file".
6

7 **Serverless Distributed File System**

8 Fig. 1 illustrates an exemplary network environment 100 that supports a
9 serverless distributed file system. Four client computing devices 102, 104, 106,
10 and 108 are coupled together via a data communications network 110. Although
11 four computing devices are illustrated, different numbers (either greater or fewer
12 than four) may be included in network environment 100.

13 Network 110 represents any of a wide variety of data communications
14 networks. Network 110 may include public portions (e.g., the Internet) as well as
15 private portions (e.g., an internal corporate Local Area Network (LAN)), as well
16 as combinations of public and private portions. Network 110 may be implemented
17 using any one or more of a wide variety of conventional communications media
18 including both wired and wireless media. Any of a wide variety of
19 communications protocols can be used to communicate data via network 110,
20 including both public and proprietary protocols. Examples of such protocols
21 include TCP/IP, IPX/SPX, NetBEUI, etc.

22 Computing devices 102-108 represent any of a wide range of computing
23 devices, and each device may be the same or different. By way of example,
24 devices 102-108 may be desktop computers, laptop computers, handheld or pocket
25

1 computers, personal digital assistants (PDAs), cellular phones, Internet appliances,
2 consumer electronics devices, gaming consoles, and so forth.

3 Two or more of devices 102-108 operate to implement a serverless
4 distributed file system. The actual devices included in the serverless distributed
5 file system can change over time, allowing new devices to be added to the system
6 and other devices to be removed from the system. Each device 102-108 that is
7 part of the distributed file system has portions of its mass storage device(s) (e.g.,
8 hard disk drive) allocated for use as either local storage or distributed storage. The
9 local storage is used for data that the user desires to store on his or her local
10 machine and not in the distributed file system structure. The distributed storage
11 portion is used for data that the user of the device (or another device) desires to
12 store within the distributed file system structure.

13 In the illustrated example of Fig. 1, certain devices connected to network
14 110 have one or more mass storage devices that include both a distributed portion
15 and a local portion. The amount allocated to distributed or local storage varies
16 among the devices. For example, device 102 has a larger percentage allocated for
17 a distributed system portion 120 in comparison to the local portion 122; device
18 104 includes a distributed system portion 124 that is approximately the same size
19 as the local portion 126; and device 106 has a smaller percentage allocated for a
20 distributed system portion 128 in comparison to the local portion 130. The storage
21 separation into multiple portions may occur on a per storage device basis (e.g., one
22 hard drive is designated for use in the distributed system while another is
23 designated solely for local use), and/or within a single storage device (e.g., part of
24 one hard drive may be designated for use in the distributed system while another
25 part is designated for local use). The amount allocated to distributed or local

1 storage may vary over time. Other devices connected to network 110, such as
2 computing device 108, may not be part of the distributed file system and thus do
3 not have any of their mass storage device(s) allocated for use by the distributed
4 system. Hence, device 108 has only a local portion 132.

5 A distributed file system 150 operates to store one or more copies of files
6 on different computing devices 102-106. When a new file is created by the user of
7 a computer, he or she has the option of storing the file on the local portion of his
8 or her computing device, or alternatively in the distributed file system. If the file
9 is stored in the distributed file system 150, the file will be stored in the distributed
10 system portion of the mass storage device(s) of one or more of devices 102-106.
11 The user creating the file typically has no ability to control which device 102-106
12 the file is stored on, nor any knowledge of which device 102-106 the file is stored
13 on. Additionally, replicated copies of the file will typically be saved, allowing the
14 user to subsequently retrieve the file even if one of the computing devices 102-106
15 on which the file is saved is unavailable (e.g., is powered-down, is malfunctioning,
16 etc.).

17 The distributed file system 150 is implemented by one or more components
18 on each of the devices 102-106, thereby obviating the need for any centralized
19 server to coordinate the file system. These components operate to determine
20 where particular files are stored, how many copies of the files are created for
21 storage on different devices, and so forth. Exactly which device will store which
22 files depends on numerous factors, including the number of devices in the
23 distributed file system, the storage space allocated to the file system from each of
24 the devices, how many copies of the file are to be saved, a cryptographically
25 secure random number, the number of files already stored on the devices, and so

1 on. Thus, the distributed file system allows the user to create and access files (as
2 well as folders or directories) without any knowledge of exactly which other
3 computing device(s) the file is being stored on.

4 The files stored by the file system are distributed among the various devices
5 102-106 and stored in encrypted form. When a new file is created, the device on
6 which the file is being created encrypts the file prior to communicating the file to
7 other device(s) for storage. The directory entry (e.g., the file name) for a new file
8 is also communicated to the other device(s) for storage. Additionally, if a new
9 folder or directory is created, the directory entry (e.g., folder name or directory
10 name) is also communicated to the other device(s) for storage. As used herein, a
11 directory entry refers to any entry that can be added to a file system directory,
12 including both file names and directory (or folder) names.

13 The distributed file system 150 is designed to prevent unauthorized users
14 from reading data stored on one of the devices 102-106. Thus, a file created by
15 device 102 and stored on device 104 is not readable by the user of device 104
16 (unless he or she is authorized to do so). In order to implement such security, the
17 contents of files as well as all directory entries are encrypted, and only authorized
18 users are given the decryption key. Thus, although device 104 may store a file
19 created by device 102, if the user of device 104 is not an authorized user of the
20 file, the user of device 104 cannot decrypt (and thus cannot read) either the
21 contents of the file or its directory entry (e.g., filename).

22 23 **File Encryption**

24 The files are encrypted using a technology known as “convergent
25 encryption”. Convergent encryption has the following two properties. First, if

1 two or more encryptable objects are identical, then even if different encryption
2 keys are used to encrypt them to provide individual cipher objects, one does not
3 need to have access to any of the encryption keys to determine from an
4 examination of the cipher objects that the encryptable objects are identical.
5 Second, if two or more encryptable objects are identical but are encrypted with
6 different encryption keys, the total space that is required to store all of the cipher
7 objects is proportional to the space that is required to store a single encryptable
8 object, plus a constant amount of storage for each distinct encryption key.

9 Generally, according to convergent encryption, a file F (or any other type of
10 encryptable object) is initially hashed using a one-way hashing function h (e.g.,
11 SHA, MD5, etc.) to produce a hash value $h(F)$. The file F is then encrypted using
12 a symmetric cipher (e.g., RC4, RC2, etc.) with the hash value as the key, or
13 $E_{h(F)}(F)$. Next, read access control entries are created for each authorized user who
14 is granted read access to the encrypted file. Write access control is governed by
15 the directory server that stores the directory entry for the file, and it is thus not
16 addressed by the file format and is not discussed further within this document. All
17 references to "access" within this document refer to read access. The access
18 control entries are formed by encrypting the file's hash value $h(F)$ with any
19 number of keys K_1, K_2, \dots, K_m , to yield $E_{K_1}(h(F)), E_{K_2}(h(F)), \dots, E_{K_m}(h(F))$. In
20 one implementation, each key K is the user's public key of a public/private key
21 pair for an asymmetric cipher (e.g., RSA).

22 With convergent encryption, one encrypted version of the file is stored and
23 replicated among the serverless distributed file system 150. Along with the
24 encrypted version of the file is stored one or more access control entries depending
25

1 upon the number of authorized users who have access. Thus, a file in the
2 distributed file system 150 has the following structure:

$$3$$
$$4 [E_{h(F)}(F), \langle E_{K1}(h(F)) \rangle, \langle E_{K2}(h(F)) \rangle, \dots, \langle E_{Km}(h(F)) \rangle]$$
$$5$$

6 One advantage of convergent encryption is that the encrypted file can be
7 evaluated by the file system to determine whether it is identical to another file
8 without resorting to any decryption (and hence, without knowledge of any
9 encryption keys). Unwanted duplicative files can be removed by adding the
10 authorized user(s) access control entries to the remaining file. Another advantage
11 is that the access control entries are very small in size, on the order of bytes as
12 compared to possibly gigabytes for the encrypted file. As a result, the amount of
13 overhead information that is stored in each file is small. This enables the property
14 that the total space used to store the file is proportional to the space that is required
15 to store a single encrypted file, plus a constant amount of storage for each
16 additional authorized reader of the file.

17 For more information on convergent encryption, the reader is directed to
18 co-pending U.S. Patent Application Serial No. 09/565,821, entitled "Encryption
19 Systems and Methods for Identifying and Coalescing Identical Objects Encrypted
20 with Different Keys", which was filed May 5, 2000, in the names of Douceur et
21 al., and is commonly assigned to Microsoft Corporation. This application is
22 hereby incorporated by reference.

23
24
25

Computing Device Architecture

Fig. 2 illustrates logical components of an exemplary computing device 200 that is representative of any one of the devices 102-106 of Fig. 1 that participate in the distributed file system 150. Computing device 200 includes a server component 202, a client component 204, a memory 206, a mass storage device 208, and a distributed file system interface 210. Computing device 200 also typically includes additional components (e.g., a processor), however these additional components have not been shown in Fig. 2 so as not to clutter the drawings. A more general description of a computer architecture with various hardware and software components is described below with reference to Fig. 3.

Memory 206 can be any of a wide variety of conventional volatile and/or nonvolatile memories, such as RAM, ROM, Flash memory, and so on. Mass storage device 208 can be any of a wide variety of conventional nonvolatile storage devices, such as a magnetic disk, optical disk, Flash memory, and so forth. Mass storage device 208 is partitioned into a distributed system portion and a local portion.

Computing device 200 is intended to be used in a serverless distributed file system, and as such includes both a server component 202 and client component 204. Server component 202 handles requests when device 200 is responding to a request involving a file or directory entry stored (or to be stored) in storage device 208, while client component 204 handles the issuance of requests by device 200 for files stored (or to be stored) in the distributed file system. Client component 204 and server component 202 operate independent of one another. Thus, situations can arise where the serverless distributed file system 150 causes files

1 being stored by client component 204 to be stored in mass storage device 208 by
2 server component 202.

3 Client component 204 includes a storage and retrieval control module 220,
4 which along with interface 210, manages access to the serverless distributed file
5 system 150 for the creation, storage, retrieval, reading, writing, modifying, and
6 verifying of files and directories on behalf of computing device 150. The control
7 module 220 uses a segmenting module 222, a cryptographic engine 224, a hashing
8 module 226, a signing/verification module 228, and tree builder 230 when
9 handling the encrypted files 240 stored in the distributed system portion of the
10 mass storage 208. These components 222-230 perform the various operations of
11 the convergent encryption process to create and maintain files, as well as
12 facilitating verification of the contents of the files without decryption. These
13 components are described in more detail below.

14 The server component 202 includes a distributed system control module
15 250 and a duplication identifier 252. Distributed system control module 250
16 manages access to the encrypted files 240. It communicates with mass storage
17 device 208 to store and retrieve encrypted files 240. Distributed system control
18 module 250 also maintains a record of the encrypted directory entries (not shown)
19 in memory 206 and/or mass storage device 208 that are stored at computing device
20 200 (or alternatively that are stored elsewhere in the serverless distributed file
21 system).

22 Duplication identifier 252 helps identify identical encrypted files in the
23 distributed file system. When the duplication identifier 252 finds a duplication
24 that is not an intentional replication for fault tolerant purposes, the duplication
25 identifier 252 notifies the control module 250, which then eliminates the

1 duplicated file and adds the access control entries to the eliminated file to the
2 remaining file.

3 Fig. 3 illustrates a more general computer environment 300, which is used
4 to implement the distributed file system. The computer environment 300 is only
5 one example of a computing environment and is not intended to suggest any
6 limitation as to the scope of use or functionality of the computer and network
7 architectures. Neither should the computer environment 300 be interpreted as
8 having any dependency or requirement relating to any one or combination of
9 components illustrated in the exemplary computer environment 300.

10 Computer environment 300 includes a general-purpose computing device in
11 the form of a computer 302. The components of computer 302 can include, by are
12 not limited to, one or more processors or processing units 304, a system memory
13 306, and a system bus 308 that couples various system components including the
14 processor 304 to the system memory 306.

15 The system bus 308 represents one or more of any of several types of bus
16 structures, including a memory bus or memory controller, a peripheral bus, an
17 accelerated graphics port, and a processor or local bus using any of a variety of
18 bus architectures. By way of example, such architectures can include an Industry
19 Standard Architecture (ISA) bus, a Micro Channel Architecture (MCA) bus, an
20 Enhanced ISA (EISA) bus, a Video Electronics Standards Association (VESA)
21 local bus, and a Peripheral Component Interconnects (PCI) bus also known as a
22 Mezzanine bus.

23 Computer 302 typically includes a variety of computer readable media.
24 Such media can be any available media that is accessible by computer 302 and
25

1 includes both volatile and non-volatile media, removable and non-removable
2 media.

3 The system memory 306 includes computer readable media in the form of
4 volatile memory, such as random access memory (RAM) 310, and/or non-volatile
5 memory, such as read only memory (ROM) 312. A basic input/output system
6 (BIOS) 314, containing the basic routines that help to transfer information
7 between elements within computer 302, such as during start-up, is stored in ROM
8 312. RAM 310 typically contains data and/or program modules that are
9 immediately accessible to and/or presently operated on by the processing unit 304.

10 Computer 302 may also include other removable/non-removable,
11 volatile/non-volatile computer storage media. By way of example, Fig. 3
12 illustrates a hard disk drive 316 for reading from and writing to a non-removable,
13 non-volatile magnetic media (not shown), a magnetic disk drive 318 for reading
14 from and writing to a removable, non-volatile magnetic disk 320 (e.g., a "floppy
15 disk"), and an optical disk drive 322 for reading from and/or writing to a
16 removable, non-volatile optical disk 324 such as a CD-ROM, DVD-ROM, or other
17 optical media. The hard disk drive 316, magnetic disk drive 318, and optical disk
18 drive 322 are each connected to the system bus 308 by one or more data media
19 interfaces 326. Alternatively, the hard disk drive 316, magnetic disk drive 318,
20 and optical disk drive 322 can be connected to the system bus 308 by one or more
21 interfaces (not shown).

22 The disk drives and their associated computer-readable media provide non-
23 volatile storage of computer readable instructions, data structures, program
24 modules, and other data for computer 302. Although the example illustrates a hard
25 disk 316, a removable magnetic disk 320, and a removable optical disk 324, it is to

1 be appreciated that other types of computer readable media which can store data
2 that is accessible by a computer, such as magnetic cassettes or other magnetic
3 storage devices, flash memory cards, CD-ROM, digital versatile disks (DVD) or
4 other optical storage, random access memories (RAM), read only memories
5 (ROM), electrically erasable programmable read-only memory (EEPROM), and
6 the like, can also be utilized to implement the exemplary computing system and
7 environment.

8 Any number of program modules can be stored on the hard disk 316,
9 magnetic disk 320, optical disk 324, ROM 312, and/or RAM 310, including by
10 way of example, an operating system 326, one or more application programs 328,
11 other program modules 330, and program data 332. Each of such operating
12 system 326, one or more application programs 328, other program modules 330,
13 and program data 332 (or some combination thereof) may implement all or part of
14 the resident components that support the distributed file system.

15 A user can enter commands and information into computer 302 via input
16 devices such as a keyboard 334 and a pointing device 336 (e.g., a "mouse").
17 Other input devices 338 (not shown specifically) may include a microphone,
18 joystick, game pad, satellite dish, serial port, scanner, and/or the like. These and
19 other input devices are connected to the processing unit 304 via input/output
20 interfaces 340 that are coupled to the system bus 308, but may be connected by
21 other interface and bus structures, such as a parallel port, game port, or a universal
22 serial bus (USB).

23 A monitor 342 or other type of display device can also be connected to the
24 system bus 308 via an interface, such as a video adapter 344. In addition to the
25 monitor 342, other output peripheral devices can include components such as

1 speakers (not shown) and a printer 346 which can be connected to computer 302
2 via the input/output interfaces 340.

3 Computer 302 can operate in a networked environment using logical
4 connections to one or more remote computers, such as a remote computing device
5 348. By way of example, the remote computing device 348 can be a personal
6 computer, portable computer, a server, a router, a network computer, a peer device
7 or other common network node, and the like. The remote computing device 348 is
8 illustrated as a portable computer that can include many or all of the elements and
9 features described herein relative to computer 302.

10 Logical connections between computer 302 and the remote computer 348
11 are depicted as a local area network (LAN) 350 and a general wide area network
12 (WAN) 352. Such networking environments are commonplace in offices,
13 enterprise-wide computer networks, intranets, and the Internet.

14 When implemented in a LAN networking environment, the computer 302 is
15 connected to a local network 350 via a network interface or adapter 354. When
16 implemented in a WAN networking environment, the computer 302 typically
17 includes a modem 356 or other means for establishing communications over the
18 wide network 352. The modem 356, which can be internal or external to computer
19 302, can be connected to the system bus 308 via the input/output interfaces 340 or
20 other appropriate mechanisms. It is to be appreciated that the illustrated network
21 connections are exemplary and that other means of establishing communication
22 link(s) between the computers 302 and 348 can be employed.

23 In a networked environment, such as that illustrated with computing
24 environment 300, program modules depicted relative to the computer 302, or
25 portions thereof, may be stored in a remote memory storage device. By way of

1 example, remote application programs 358 reside on a memory device of remote
2 computer 348. For purposes of illustration, application programs and other
3 executable program components such as the operating system are illustrated herein
4 as discrete blocks, although it is recognized that such programs and components
5 reside at various times in different storage components of the computing device
6 302, and are executed by the data processor(s) of the computer.

7 An implementation of the distributed file system 150 may be described in
8 the general context of computer-executable instructions, such as program modules,
9 executed by one or more computers or other devices. Generally, program modules
10 include routines, programs, objects, components, data structures, etc. that perform
11 particular tasks or implement particular abstract data types. Typically, the
12 functionality of the program modules may be combined or distributed as desired in
13 various embodiments.

14 An implementation of the file format for the encrypted files may be stored
15 on or transmitted across some form of computer readable media. Computer
16 readable media can be any available media that can be accessed by a computer.
17 By way of example, and not limitation, computer readable media may comprise
18 "computer storage media" and "communications media."

19 "Computer storage media" includes volatile and non-volatile, removable
20 and non-removable media implemented in any method or technology for storage
21 of information such as computer readable instructions, data structures, program
22 modules, or other data. Computer storage media includes, but is not limited to,
23 RAM, ROM, EEPROM, flash memory or other memory technology, CD-ROM,
24 digital versatile disks (DVD) or other optical storage, magnetic cassettes, magnetic
25 tape, magnetic disk storage or other magnetic storage devices, or any other

1 medium which can be used to store the desired information and which can be
2 accessed by a computer.

3 "Communication media" typically embodies computer readable
4 instructions, data structures, program modules, or other data in a modulated data
5 signal, such as carrier wave or other transport mechanism. Communication media
6 also includes any information delivery media. The term "modulated data signal"
7 means a signal that has one or more of its characteristics set or changed in such a
8 manner as to encode information in the signal. By way of example, and not
9 limitation, communication media includes wired media such as a wired network or
10 direct-wired connection, and wireless media such as acoustic, RF, infrared, and
11 other wireless media. Combinations of any of the above are also included within
12 the scope of computer readable media.

13 **File Format**

14
15 Fig. 4 generally shows a file format 400 of a file that is stored in the
16 distributed file system 150. The file format is composed of two parts: a data
17 stream 402 and a metadata stream 404. The data stream 402 contains the contents
18 of the file, which makes up the bulk of the entire file. The data stream 402 is a
19 primary (unnamed) data stream that may be analyzed using the single instance
20 store (SIS) technology introduced by Microsoft Corporation and discussed in the
21 Background section. SIS components work on unnamed streams, while ignoring
22 other streams. The metadata stream 404 of the file 400 is a separate named stream
23 that is ignored by the SIS system.

24 For discussion purposes, each file is described as containing only a single
25 user-defined stream, and that stream is the unnamed data stream. However, the

distributed file system 150 is capable of supporting any number of user-defined streams per file. A file naming convention differentiates among multiple streams by prepending a user-defined stream name to a file system stream name "FSMetadata\$". So, if a user file has a named stream called "alternateStream", the file would have a second metadata stream called "alternateStreamFSMetadata\$." If a user has the bad taste to name a stream with a string that ends in "FSMetadata\$" followed by 0 or more "\$"s, the file system changes the name of the stream by adding an additional "\$".

Data Stream 402

The data stream 402 is designed to allow efficient verification, reading, and writing of portions of the file, without affecting other portions. The data stream is encrypted using the convergent encryption technology described above beneath the heading "File Encryption". For small files, the entire file is hashed and encrypted using the resulting hash value as the encryption key. The encrypted file can be verified without knowledge of the key or any need to decrypt the file first.

For large files, however, it is difficult to read or update only part of a file because the encryption of the file is based on a hash of the entire file contents. Any write to a file would require re-hashing the entire file followed by re-encrypting with the newly generated hash as the key. Furthermore, verification involves hashing the entire file and examining the hash value. Taking a single hash of the ciphertext of a large file for verification purposes makes writes to part of the file expensive, because any write would once again require hashing the whole file.

1 To overcome this problem for large files, the file contents in the data stream
2 402 can be broken into smaller blocks and then convergent encryption can be
3 applied separately to each block. Fig. 5 shows a file F that is segmented by
4 segmenting module 222 into an array 500 of multiple blocks 502(0)-502(n-1). In
5 one implementation, the blocks are fixed sized chunks. For example, the file F is
6 segmented into "n" pages F^0 - F^{n-1} , where each page is a fixed size. A 4Kbyte size
7 is one suitable size of each block because it is the smallest page size used by NT
8 systems (although some architectures use multiples of 4K). Hence, it is the
9 smallest chunk that the memory manager is going to request on a cache miss or
10 user mapped file page fault, and it is also the smallest chunk that is going to be
11 written by the lazy writer.

12 From the system perspective, each page is originally deemed as "cleartext",
13 meaning that it has not yet undergone encryption as part of the convergent
14 encryption process. The file F may actually be in a plain, unencrypted form, or it
15 may have already been encrypted in some manner. Thus, by noting that each file
16 page is "cleartext", we are simply explaining the process from the point of view of
17 the convergent encryption process, regardless of the condition in which file F
18 originally resides.

19 Convergent encryption is then applied to the file at the block level. That is,
20 each block F^i is separately hashed using a one-way hash function (e.g., SHA,
21 MD5, etc.) to produce a hash value $h(F^i)$. Each block F^i is then encrypted using a
22 symmetric cipher (e.g., RC4, RC2, etc.) and the hash value $h(F^i)$ as the key, or
23 $E_{h(F^i)}(F^i)$. This produces an array 504 of encrypted blocks 506(0)-506(n-1).

24 The encrypted blocks 506 form the contents of the unnamed data stream
25 402 in file 400. That is, the data stream 402 contains encrypted blocks $E_{h(F^i)}(F^i)$

for all i from 0 to the size of the file in pages minus one. Nothing else goes in the main stream.

Metadata Stream 404

Returning to Fig. 4, the metadata stream 404 of the file 400 stores metadata used to describe the contents of the file and to decrypt the file. The metadata stream 404 contains a header 406, a tree structure 408, and some per user information 410.

Header 406

The header 406 contains information pertaining to the file and which may be used to validate the file. In Fig. 4, the header 406 is illustrated as including such file information as a file number 412, a revision number 414, a hash value 416 of the root of tree structure 408, and an optional digital signature 418.

Exemplary implementations of the header will be described according to three different file formats, each of which possess different advantages. The three formats include a signed format, an unsigned format, and a delegation format. The signed format contains a digital signature 418 associated with the file and information used to verify the signature. The unsigned file format omits the digital signature and verification information. The delegation format accepts delegation certificates to convey ownership or privileges with respect to the file. The signed and unsigned formats are described immediately below, while the delegation format is described further along in this discussion beneath the heading "File Format Using Delegation Certificates".

With the unsigned format, the directory servers send a hash value that represents the contents of the file (as well as its metadata) to a verifying machine for verification of the file contents. In this implementation, there is no way for the verifying machine to determine that a particular user wrote a file, aside from trusting the directory servers. The advantage of this approach is that there is no need to compute or verify digital signatures for the file, which can be computationally expensive. The disadvantage is that the verifying machine must trust the directory servers, although this is less of a problem than it might seem on its face. Even with the signed format, where it is possible to verify that a particular user wrote a file without trusting the directory servers, one still relies on the directory servers to verify that the correct version of the file is present (as opposed to a different file or a old version of the correct file), and to state which user(s) are allowed to sign a particular file. In practice, corrupt directory servers could do much damage even with signed files, so electing to use unsigned files and dropping the signatures saves computational cost at a slight increased risk of reliance on the directory servers.

The following example header contains fields common to all three formats. Fields marked with (SF) exist only in the signed format that uses signatures.

ULONG	MagicNumber;	(must be 0x0fa2317e)
UCHAR	FormatVersionMajor;	(1 described here)
UCHAR	FormatVersionMinor;	(1 described here)
USHORT	Flags;	1 means alternate signature type 2 means short header format
GUID	FileId;	
FILE_USER_NAME	FileOwner;	
Time	FileCreateTime;	E.g., 64 bit NT time
Time	FileModifyTime;	
LONGLONG	FileRevisionNumber;	(SF)
FILE_USER_NAME	LastWriter;	
LONGLONG	FileSize;	
LONGLONG	FragmentOffset;	
LONGLONG	FragmentSize;	
ULONG	UserKeyCount;	The number of KEY_ID_PAIRS


```

        ULONG      FilePageSize;
        ULONG      BytesPerTreeBlock;
        USHORT     PublicKeyAlgorithm; 1 is RSA, others undefined
        USHORT     HashAlgorithm;      1 is SHA, 2 is MD5
        USHORT     SymmetricAlgorithm; 1 is RC4, 2 is RC2
        USHORT     SymmetricKeySize;
        USHORT     SignatureSize;      (SF)
        LONGLONG   KeyPairOffset;      File offset of the KEY_ID_PAIRs
        LONGLONG   KeyDataOffset;
        LONGLONG   SignatureOffset;    (SF) The file offset of the sig

```

```

@ SignatureOffset: (SF)
    <Signature List, signed by the LastWriter, size SignatureSize>

```

```

@ KeyPairOffset:
    KEY_ID_PAIR    UserKeys[UserKeyCount];

```

```

@ KeyDataOffset:
    The key data, as described below.

```

The MagicNumber field characterizes the type of header. Two FormatVersion fields describe the format version of the encrypted file itself and are intended to be used to allow an on-disk format to evolve over time. If the distributed file system 150 encounters a newer FormatVersionMajor than it presently understands, the file system is unable to understand the file. The file system ignores FormatVersionMinors that are too big and treats the format as if it were the newest understood by the file system and consistent with the FormatVersionMajor.

The Flags field contains a first flag to differentiate between whether the signed format or the delegation format is employed. The Flag field also contains a second flag to differentiate between a short header form and a long header form. If the second flag is set to indicate a short header form, the rest of the header after the FileModifyTime field uses the following format rather than the one presented above:

```

        USHORT     FileRevisionNumber; (SF)
        USHORT     FileSize;

```

```

USHORT          SignatureSize;      (SF)
1
<Signature List, signed by the FileOwner, size SignatureSize> (SF)
2
<A single KEY_ID_PAIR goes here>
3

```

4 The short header form is intended for small files (typically 4K or less,
5 although they can be as big as 64K). The conditions in which the short header
6 form can be used are:

- 8 • The LastWriter and FileOwner are the same;
- 9 • There is only one key-id pair entry;
- 10 • The file uses RSA/SHA and constant size (e.g., 128-bit) symmetric
11 keys;
- 12 • The revision number fits in a USHORT in the signed format case;
13 and
- 14 • Employs a single file encryption/cleartext hash (as is done with files
15 that are less than or equal to one file page size in the normal format).

16
17 The short header form is intended for the extremely common case of very
18 small files that are created once and rarely (or never) overwritten, and readable by
19 either everyone or just the creator. This may be as many as half of all files. The
20 short header form is incompatible with the delegation format, but since these files
21 are created in one piece and then left alone and the delegation format is intended
22 to address in-place updates, the incompatibility is not a problem. The distributed
23 file system is free to decide whether to use the short or long header form for any
24 particular file, and can switch formats on the file (assuming that it has access to
25 the writing user's key).

1 The FileId field contains the file number in the form of a globally unique
2 ID. The FileOwner field identifies the file owner, the FileCreateTime field
3 specifies the time of file creation, and the FileModifyTime field specifies the last
4 time the file was modified. The FileRevisionNumber field, which is only present
5 in the signed format, is updated every time a file is written, closed, and signed
6 (i.e., not for every write to the file). The directory servers will know what the
7 latest revision of a file is.

8 The LastWriter field notes the last user to write to the file. This user is also
9 the one who issued the signature for the file contents stored at the offset held in the
10 SignatureOffset field. The type FILE_USER_NAME that is used for the
11 LastWriter and FileOwner fields supports two globally unique identifiers, one for
12 the user and one for an authority that certifies the identity of the user, although
13 there are a number of other possibilities for user names, including a hash of the
14 user's public key. The FileSize field describes the size of the entire file, while the
15 FilePageSize field specifies the size of each page in the file.

16 The FragmentOffset and FragmentSize fields are intended to support very
17 large files that have been broken into fragments to make them more manageable
18 by the replica placement and regeneration systems. As one example default, the
19 FragmentOffset field is set to zero and the FragmentSize field equals FileSize.

20 The PublicKeyAlgorithm field specifies a suitable public key cipher, such
21 as RSA. The HashAlgorithm field identifies a suitable hash algorithm, such as
22 SHA or MD5. The SymmetricAlgorithm field specifies a suitable symmetric
23 cipher, such as RC2 or RC4, and it employs keys of a size specified in the
24 SymmetricKeySize field (e.g., 128 bit).

1 The KeyPairOffset field contains an offset value to a location in the
2 metadata stream that holds the key ID pairs (per user information 410). The
3 KeyDataOffset field contains an offset value to a location in the metadata stream
4 where the tree 408 is located.

6 Tree Structure 408

7 The tree 408 is the portion of the metadata stream 404 that facilitates
8 indexing into individual blocks in the data stream 402, thereby enabling data
9 verification of the contents in the data stream. The tree contains data for two
10 purposes: (1) allow a user to decrypt the file one block at a time and out of order,
11 and (2) allow data servers to verify that the contents of the file is genuine one
12 block at a time and out of order and without having access to the keys of any of
13 the authorized readers of the file. In this manner, the tree 408 allows the
14 distributed file system 150 to verify individual encrypted blocks 506 directly,
15 without decryption and without any knowledge of the encryption keys used to
16 encrypt the file.

17 Fig. 6 shows a tree structure 408 in more detail. There are two types of
18 entries in tree 408: leaf entries and higher-order entries. They differ in that the
19 decryption keys are held only in the leaf entries, while both the leaf entries and the
20 higher-order entries contain hashes that are used to determine whether the file
21 contents are correct.

22 In Fig. 6, the tree 408 defines leaf nodes 602(0)-602(n-1) for each of the
23 corresponding encrypted blocks 506(0)-506(n-1). Each leaf node L^i contains two
24 components: (1) an access value 604 used for decrypting the corresponding block
25 and (2) a verification value 606 used for verifying the corresponding block. In the

illustrated implementation, the access value 604 is formed by encrypting a hash of the cleartext file block using a symmetric cipher E and a randomly generated key K, or

$$\text{Access Value} = E_K(h(F^i)).$$

The symmetric cipher specified in the SymmetricAlgorithm field of the header is used for this encryption (e.g., RC2 or RC4).

The verification value 606 is created by hashing the associated encrypted block, or:

$$\text{Verification value} = h(E_{h(F^i)}(F^i)).$$

A leaf entry possesses the following format:

BYTE	EncryptedCleartextHash[HASH_SIZE];
BYTE	UnencryptedCiphertextHash[HASH_SIZE];

where EncryptedCleartextHash is the access value 604 and the UnencryptedCiphertextHash is the verification value 606. The HASH_SIZE value depends on the HashAlgorithm specified in the header 406. For the SHA algorithm, it is 20 bytes and for the MD5 algorithm, it is 16 bytes. When RC2 encryption is used for the EncryptedCleartextHash, the size is 24 bytes due to padding, regardless of which hash algorithm (MD5 or SHA) is used.

The existence and size of the tree 408 varies with the size of the file. At one extreme, if the file is less than or equal to one page in size, there is no tree and

1 no per-file secret key K. Instead, the cleartext hash value is turned into a key,
2 encrypted with the public keys of the readers, and stored in the
3 FILE_KEY_ID_PAIR. The hash of the file (that is stored at the directory servers
4 or signed and placed in the file) includes the entire contents of the ciphertext of the
5 file in place of the hash of the highest level that exists in the tree. Since about half
6 of all files are smaller than 4K, this optimization can be significant.

7 For slightly larger files, the tree may be one level deep, containing only the
8 leaf nodes 602 and one root node formed from the leaf nodes. As a general rule, if
9 the count of pages in a file is greater than one but less than a moderate value
10 obtained by dividing the BytesPerTreeBlock by the leaf entry size (i.e.,
11 approximately 3.2 Mbytes for SHA, 4K file pages and BytesPerTreeBlock of 32
12 Kbytes), the tree 408 only contains leaf nodes and no higher order entries. The
13 number of leaf nodes 602 is sufficient to contain enough entries to describe all of
14 the pages in the file. For the signed format case, the LastWriter signs the header,
15 the per user information, and one or more verification values from the tree, as
16 described in more detail below. In the unsigned format case, the directory servers
17 store the hash that the LastWriter would have signed.

18 For large files, the cost of computing the hashes can be quite high. For
19 example, for a 500Mbyte file (e.g., an email file) and 4K file pages, there are
20 approximately 128,000 hashes. At twenty bytes per hash, a single block update to
21 this file would require 2.5MB of hashing.

22 To reduce the amount of work for small writes to large files, the tree may
23 be configured with one or more intermediate levels of hashes. In Fig. 6, the leaf
24 nodes 602 are grouped into tree blocks (e.g., denoted as TB_0^0 to mean tree block 0
25 at tree level 0). The size of a block is specified in the BytesPerTreeBlock field in

header 406. If the block size does not divide evenly by the size of a leaf (or higher order) entry, the block is padded with zeroes after the last complete entry.

The tree block is hashed using a one-way hashing function to form intermediate nodes 610(0)-610(j). A higher order intermediate node has the following format:

```
BYTE                                     HashOfLowerOrderEntry[HASH_SIZE];
```

Each hash is of an entire block of the next lower level of the tree, excluding any trailing padding and excluding the unused portion of the final leaf block. The first higher order block follows the first complete leaf tree block. Unlike the leaf tree blocks, the higher order blocks are allocated as a whole because if there is a higher order block, there is also a leaf tree block following it. By allocating the whole block, the system avoids having to move the whole structure around to expand the file. Using the example parameters from above, a first higher order block is not used until the file is approximately 3.2 MB. Thus, the maximum wasted space for a higher order block is 1% (32Kb / 3.2MB), which is on the order of the cost of the leaf tree entries.

The array of intermediate nodes 610 can again be grouped into blocks (e.g., denoted as TB_1^0 to mean tree block 0 at tree level 1) and each block is hashed to form the next level of nodes represented by node 620(0). The grouping and hashing process can be repeated as many times as desired until reaching a tree root, which is denoted as R_x^0 . After the first higher order block follows more leaf tree blocks until all of the entries in the higher order block are filled, at which time follows another leaf block and the next higher order block, followed by the rest of

the leaf blocks for the second higher order block, the first leaf block for the third higher order block, the third higher order block, and so on. It is noted, however, that the second higher order block is rarely used since it typically is introduced for very large files of approximately 5.2GB or greater (using the example parameters).

The layout has the following recursive definition, in which the superscripts are eliminated for clarity:

$$C_0 = L$$

$$C_x = C_{x-1}H_xC_{x-1}C_{x-1}\dots C_{x-1}$$

where H_x indicates an x^{th} -order higher-level block. So, the layout of the data in a two-level tree is as follows:

$$L^0H_1^0L^1L^2\dots L^{n-1}H_2^0L^nH_1^1L^{n+1}\dots L^{2n-1}L^{2n}H_1^2L^{2n+1}\dots$$

where n is the number of entries in the higher order block.

The tree root is then hashed to form the root node 630, or $h(R_x^0)$. This hash value may then be hashed together with the metadata header 406 and per user information 410 and the resulting hash stored at the directory servers in the case of the unsigned format, or signed using a user's signature in the signed format case. In this way, the hash or signature covers the higher order blocks (of the highest order that exists in the file) and thereby indirectly covers the leaf blocks. The signature covers the used entries in the higher order blocks, not the unused entries and padding. Similarly, the hash entry in the higher order block of the final leaf block does not include any unused entries/padding in that leaf block.

1 With this tree structure, every small update to the file merely involves
2 changing the file block, the leaf node associated with the file block, and the nodes
3 in the tree branch to the leaf node. This solution reduces the hashing cost because
4 the number of upper-level hashes that need to be modified for any given write
5 grows logarithmically in the size of the file. Therefore, with the hash tree, any
6 work to update a particular byte of the file is proportional to the depth of the tree,
7 which grows with the log of the size of the file.

8 It is noted that although a multi-level tree index is described herein as one
9 possible implementation, other forms of indexing structures may be used.

11 User Key List

12 To grant access privileges to multiple users, the file system 150 maintains a
13 user key list for each file. Each entry in the user key list contains the data used by
14 specific users to decrypt the file. More particularly, the randomly generated key
15 K, which is used to encrypt the hash of the blocks of the cleartext file F (i.e.,
16 forming the access value 604), is encrypted using each authorized user's public
17 key and stored in a user key list, or $E_{\text{UiPubKey}}(K)$.

18 Fig. 7 illustrates a user key list 700 for the file F. Each entry in list 700
19 includes a user name 702 of the user with access privileges and the encrypted
20 symmetric key 704. If the file is publicly readable, the key list 700 contains only a
21 special entry 706, in which the FILE_USER_NAME is the reserved value
22 USER_EVERYONE to indicate that everyone has access, and the associated
23 encrypted key segment contains the key data in the clear.

24 Each entry in the key list 700 has the following format:
25

FARSITE_USER_NAME	UserName;
ULONG	EncryptedKeyBlobSize;
BYTE	EncryptedKeyBlob[EncryptedKeyBlobSize];

The content of the EncryptedKeyBlob field depends on the size of the file. If the file is one page or less in size, the field contains the key derived from the hash of the cleartext of the file, encrypted with the public key of UserName. If the file is bigger than one page, the field contains the random symmetric key K that was used to encrypt the hashes of the cleartext of the file in the leaf tree block(s), also encrypted with the public key of UserName.

There are two different types of signatures for a file, depending on how the file is written. In the signed format, the file is signed by the user who is named in the LastWriter field. The signature covers the file header (up to and including the SignatureSize), but does not cover the three offsets so that servers can re-arrange pieces of the Metadata\$ stream as they see fit, without having the last writer's key. After the header, the signature then covers the key-id pairs. Following that, it covers either the file ciphertext, the single leaf tree block, or the highest order tree block, depending on the file's size. In the signed format, all that is stored at SignatureOffset is the actual signature blob.

For small files, the user key list 700 does not contain entries with encrypted symmetric keys. Instead, each entry contains a user name (i.e., FILE_USER_NAME 702) and an encrypted version of the hash of the entire file (i.e., $h(F)$), which is encrypted using the user's public key. Accordingly, this portion of the entry would resemble $E_{UiPubKey}(h(F))$.

Since files are stored on machines that are not trusted, read access cannot be sufficiently controlled merely by listing the authorized readers of a file in the metadata, as is commonly done in trusted file systems. Thus, this file format relies

1 on cryptography to provide access security. Only a truly authorized user with
2 knowledge of an appropriate private key will be able to recover the access key K.
3 As a result, an impostor who attempts to recover the key K using an authorized
4 user's name will be unable to decrypt the access key K because that impostor does
5 not have knowledge of the user's private key.

6 An alternative technique may be used in the case where the user creating
7 the file (i.e., the user who makes up K) is the same as the user in the user key list
8 702. In this case, a secret symmetric key that is known only by that user can be
9 used in place of that user's public key $U_{iPubKey}$. Since symmetric key operations
10 are substantially cheaper than public key operations in terms of computational
11 resources, creating and reading the file would be computationally cheaper in the
12 common case that the file creator is the same as the file reader.

13 File Format Using Delegation Certificates

14 The third type of file format (in addition to the signed and unsigned
15 formats) is one in which delegation certificates are used in place of digital
16 signatures. Setting the first flag in the Flags field of the header 406 signifies the
17 delegation format. The delegation format is used to handle a case where a
18 machine crashed while in the process of writing a file, before the last writer signed
19 that file. With this format, a user's machine may create a delegation certificate
20 allowing other entities to verify as a group the authenticity of the file on behalf of
21 the user in the event the user's machine is unavailable to make the verification.
22

23 When a computing device attempts to write a file to the distributed file
24 system 150 and receives a write lock for a file or directory, the computing device
25 generates a random symmetric key, known as the "lock-secret" key. The

1 computing device uses secret sharing to break the lock-secret key into multiple
2 pieces, one piece for each of the directory servers, with a specified number of the
3 servers being sufficient to recover the key.

4 If the computing device wants to commit updates to a file without attaching
5 a full signature to the file (such as on a write-through write to a database file), the
6 computing device generates a delegation certificate and signs the certificate with
7 the user's private key. When the computing device updates a file, it computes the
8 hash of the file that would normally be signed with the writer's private key.
9 However, instead of signing the update, the computing device encrypts it with the
10 lock-secret key using the symmetric signature algorithm specified in the file
11 header.

12 If a machine crashes with file updates that are signed with the symmetric
13 signature key (rather than with the normal private-key signature), there will be a
14 set of files signed by lock-secret keys on recovery. For each particular lock-secret
15 key, the computing device takes all files signed by that key and sends the
16 delegation certificates and "symmetric key signatures" to all available directory
17 servers. Once the directory servers have collected all of the appropriate data, they
18 break the seal on the lock-secret key and determine whether the hash of the lock-
19 secret key matches the hash in the DelegationCertificate field. The directory
20 server then decrypt the symmetric key signature (i.e., decrypt the file hash with the
21 lock-secret key) and fill out and sign a DelegationCountersign using the decrypted
22 file hash.

23 In the delegation format, the following structure is stored at the
24 SignatureOffset field of the header:
25

```

1      LONGLONG      DelegationCertificateOffset;
2      LONGLONG      DirectoryServerSignaturesOffset;

3      @ DelegationCertificateOffset is:
4
5      ULONG          Magic;                (must be 0xdellca7e)
6      UCHAR          FormatVersionMajor; (1 described here)
7      UCHAR          FormatVersionMinor; (1 described here)
8      USHORT         HashedKeySize;
9      Time           DelegationTime;
10     GUID            FileId;
11     GUID            DelegationCertificateId;
12     LONGLONG        FileVersionNumber;
13     FILE_USER_NAME  LastWriterName;
14     ULONG            DirectoryServerCount;
15     ULONG            NumDirectoryCountersignsNeededForValidity;
16     FILE_MACHINE_NAME DirectoryServer[DirectoryServerCount];
17     ULONG            SignatureSize;
18     <a hash of the secret "signature" symmetric key, of HashedKeySize, using
19     the hash algorithm specified in the file header>
20     <the signature blob of the LastWriter >

```

```

21     @ DirectoryServerSignatureOffset is:
22     ULONG            CountOfSigningDirectoryServers;
23     for each signing server there is a DelegationCountersign:
24     ULONG            Magic                (must be
25     0xc2a38452)
26     UCHAR            FormatVersionMajor; (1 described here)
27     UCHAR            FormatVersionMinor; (1 described here)
28     USHORT           HashSize;
29     FARSITE_MACHINE_NAME SigningMachine;
30     GUID             FileId;
31     GUID             DelegationCertificateId;
32     LONGLONG          FileVersionNumber;
33     Time              CountersignTime;
34     ULONG             SignatureSize;
35     <A hash for the file contents, computed just as the hash that the
36     last writer would sign in the normal signature method, of
37     HashSize>
38     <A signature of the directory server certificate up to but not
39     including SignatureSize, followed by the file contents hash>

```

The signature in the delegation certificate covers everything from the Magic field up to but not including the SignatureSize field, and then the hash of the secret symmetric signature key. The signature of the directory servers is over what would have been signed by the user identified in the LastWriter field in the signed format case. Note that there is a separate SignatureSize for each of the directory servers, since they may have different key lengths and so different signature lengths.

1 To validate a file using the delegation format, the verifying computer first
2 evaluates the signature on the delegation certificate and confirms that the
3 certificate has the correct FileId and FileVersionNumber. It then counts the
4 number of valid DelegationCountersign's, and if that number is at least
5 NumDirectoryCountersignsNeededForValididty then the file is valid.

6 To check a DelegationCountersign, the verifying computer verifies that
7 SigningMachine is on the list in the DelegationCertificate, that the FileId,
8 FileVersionNumber and DelegationCertificateId match the DelegationCertificate,
9 and that the hash value is the same as the hash value that would have been signed
10 by the last file writer in the normal signed file format.

11 There is a related technique for the non-signature case. As before, when a
12 computing device attempts to write a file to the distributed file system 150 and
13 receives a write lock for a file or directory, the computing device generates a
14 symmetric encryption key called the "lock-secret key." The computing device
15 breaks the lock-secret key into multiple pieces and distributes the pieces to the
16 directory servers using a cryptographic secret sharing technique.

17 If the computing device subsequently wants to commit updates to a file
18 without attaching a signature, the computing device encrypts the updates with the
19 lock-secret key using the symmetric encryption algorithm specified in the file
20 header. If the directory servers are satisfied with the result, the servers accept the
21 file contents as being valid and update their internal data structures. In this case,
22 the process of producing a delegation certificate and countersigning certificate are
23 eliminated.

File Construction

Fig. 8 shows a process 800 for constructing a file according to the format shown in Figs. 4-7 for storage in the distributed file system 150. The process can be implemented in software as computer executable instructions that, when executed, perform the operations depicted in blocks. The process 800 will be described with reference to components in the computing device 200 shown in Fig. 2 as exemplary mechanisms for performing the operations, and with reference to the file format illustrated in Figs. 4-6.

The file construction process 800 builds files differently depending upon their size. Accordingly, at operation 802, a preliminary inquiry is to ascertain the size of the file. If it is a small file (e.g., 4K or less), the storage/retrieval control module 220 of the client component 204 encrypts the entire file using convergent encryption techniques (operation 804).

Alternatively, assuming the file is not small (i.e., the “no” branch from operation 802), the control module 220 proceeds to a technique for constructing a large file for storage in the distributed file system. Large file construction can be conceptualized as two phases: a first phase for creating the data stream 402 and a second phase for creating the metadata stream 404.

In the first phase, the control module 220 employs the segmenter 222 to divide a file F into “ n ” multiple blocks 502(0)-502($n-1$) at operation 806. Each block contains a portion of the file, which is illustrated as file segments F^0 , F^1 , F^2 , ..., F^{n-1} in blocks 502. At operation 808, the control module 220 invokes the hash module 226 to hash each block 502(0)-502($n-1$) to produce intermediate hash values $h(F^i)$. At operation 810, the control module 220 calls the cryptographic engine 224 to encrypt each block 502(0)-502($n-1$) using that block’s hash value, or

1 $E_{h(F^i)}(F^i)$. In practice, the hashing and encrypting operations may be accomplished
2 sequentially for each block, one block at a time, before proceeding to the next
3 block. For instance, for each block, a loop may be used to compute the hash of the
4 block, encrypt the result, and then proceed to the next block. With this approach,
5 the two accesses to the block are close together in time, which increases the
6 likelihood that the data for the block will be found in the cache and so be faster to
7 perform. The segmented and encrypted file can then be stored as the unnamed
8 data stream 402.

9 During the second phase, the control module 220 uses the tree builder
10 module 230 to construct the block-level access tree. At operation 812, the tree
11 builder 230 (or other module in the client component) generates a random K for
12 the entire file. The tree builder 230 then creates a leaf node L^i for each block
13 502(0)-502(n-1) (operation 814). Each leaf node L^i contains two components: (1)
14 an access value 604 used for decrypting the corresponding block and (2) a
15 verification value 606 used for verifying the corresponding block. Accordingly,
16 operation 814 can be viewed as two parts. At the first part represented by
17 operation 814(A), the tree builder 230 computes the access value by encrypting
18 the file segment hash $h(F^i)$ using the key K , or $E_K(h(F^i))$. At the second part
19 represented by operation 814(B), the tree builder 230 computes the verification
20 value by hashing the corresponding encrypted file segment, or $h(E_{h(F^i)}(F^i))$.

21 At operation 816, the tree builder 230 ascertains whether the tree structure
22 would benefit from an intermediate level of nodes in terms of access and
23 verification efficiency. The number of levels in the tree generally depends on the
24 size of the file and the desired fan-out. For a smaller file (e.g., a file that is greater
25 than 4KB but less than 3.2 MB), the tree is one level deep, containing only the leaf

1 nodes. For a larger file (e.g., one that is greater than 3.2 MB), another level of
2 nodes is added to enable more efficient access of the leaf nodes.

3 If another level of nodes is desired (i.e., the "yes" branch from operation
4 816), the tree builder 230 groups sets of contiguous leaf nodes to form tree blocks
5 TB_m^n (operation 818). Then, at operation 820, each tree block is hashed to form
6 intermediate tree nodes 610(0)-610(j). The process continues at operation 816,
7 where the tree builder 230 again determines whether a further level of intermediate
8 nodes would prove useful. If the file is very large (e.g., greater than 5.2GB), the
9 tree might include a second level of intermediate nodes. In this case, operations
10 818 and 820 are repeated such that the intermediate nodes in the first level are
11 grouped together to construct a second level of tree blocks (block 818) and each
12 tree block is hashed (block 820). Depending on the file size, this process is
13 repeated as many times as needed until the highest-level node contains only one
14 block.

15 Once an effective node structure is created and no more intermediate nodes
16 are desired (i.e., the "no" branch from operation 816), the tree builder 230 forms
17 the root R_x^0 and hashes it to form a hash value $h(R_x^0)$ (operation 822). In the case
18 of the signed format, the control module 220 invokes the signing/verification
19 module 228 to sign the file header 406, per-user information 410 and root node
20 $h(R_x^0)$ with the digital signature of the user identified in the LastWriter field
21 (operation 824). The resultant tree structure 408 is stored in the metadata stream
22 404. The signature is stored in the header 406 of the metadata stream 404.

File Verification

Fig. 9 shows a process 900 for verifying portions of a large file in its segmented and encrypted form, without requiring knowledge of the user private keys or random keys. For discussion of this process, it is assumed that the file is of sufficient size to have a tree structure 408 stored in the metadata stream 404. The process 900 can be implemented in software as computer executable instructions that, when executed, perform the operations depicted in blocks. The process 900 will be described with reference to components in the computing device 200 shown in Fig. 2 as exemplary mechanisms for performing the operations, and with reference to the file format illustrated in Figs. 4-6.

For discussion purposes, suppose that computing device 200 is a verifying machine that is tasked with verifying the first encrypted file block 506(0) for file segment F^0 . At operation 902, the signing/verification module 228 evaluates the signature (if any) on the header 406, per-user information 410 and tree root of the tree structure 408 using the public key of the last writer as indicated in the header 406. The signature is held in the header 406 of the metadata stream 404. If the signature is not valid (i.e., the “no” branch from operation 904), the file block is deemed not to be authentic (i.e., block 906). In the non-signed embodiment, the signing/verification module 228 computes the hash that would have been signed in the signed format case, and compares that against the has provided from the directory servers. If the hash does not match, then it follows the “no” branch from operation 904.

Conversely, if the signature is valid (i.e., the “yes” branch from operation 904), the verification module 228 verifies whether the hash value stored at the root matches the hash of the lower-order nodes below the root in the tree (i.e.,

operation 908). If the values do not match, the file block is not authentic (i.e., operation 906). If the hash is verified (i.e., the “yes” branch from operation 908), the verification module 228 traverses the tree, node by node, from the root to the leaf node L^0 associated with the target block 506(0). At operation 910, the verification module 228 moves to the next node on the path between the root and the leaf node. If the next node is not a leaf node (i.e., the “no” branch from operation 912), the verification module 228 verifies whether the hash value stored at the next node matches the hash of the lower-order nodes below that node in the tree (i.e., operation 908). In this manner, each node in the path from the root to the leaf node are evaluated. If any one of these verifications fails, the block is not authentic.

Once the leaf node is reached (i.e., the “yes” branch from block 912), at operation 916, the verification module 228 calls the hash module 226 to compute a hash of the encrypted file segment in target block, or $h(E_{h(F0)}(F^0))$. The verification module 228 then compares this resultant hash value with the verification value 606 stored in the corresponding leaf node L^0 (i.e., operation 918). If the two match (i.e., the “yes” branch from operation 920), the target block 506(0) is authentic (operation 922). If the two fail to match (i.e., the “no” branch from operation 920), the target block 506(0) is not authentic (operation 906).

Reading A File

Fig. 10 shows a process 1000 for reading one of the file blocks of a large file without having to read the entire file. As before, it is assumed that the file includes a tree structure 408 and that the target block is a block F^i . The process

1 1000 can be implemented in software and will be described with reference to
2 components in the computing device 200 in Fig. 2 and the file format in Figs. 4-7.

3 At operation 1002, the storage/retrieval control module 220 obtains the
4 random access key K from the user key list 700 by indexing into the list using the
5 File_User_Name 702 of the user who wants to read the target file block. The
6 control module 220 extracts the encrypted symmetric key 704 and decrypts the
7 access key K using the public key cipher (e.g., RSA) and the user's private key. It
8 is noted that if the user is not identified in the user key list 700, the user does not
9 have access privileges to read the file and will be prevented from doing so.
10 Additionally, an impostor attempting to recover the key K using the user's name
11 will be unable to decrypt the access key K because that impostor does not have
12 knowledge of the user's private key.

13 At operation 1004, the control module indexes into the first level of the tree
14 structure 408 in the metadata stream 404 to the leaf node L^i associated with the
15 target file block F^i . At operation 1006, the control module 220 removes the access
16 value 604 (i.e., $E_K(h(F^i))$) and calls the cryptographic engine 224 to decrypt the
17 access value using the symmetric cipher D and the symmetric access key K to
18 recover a hash of the target block, as follows:

$$D_K(E_K(h(F^i))) = h(F^i).$$

22 At operation 1008, the control module 220 calls again on the cryptographic
23 engine 224 to decrypt the target file block using a symmetric cipher D and the
24 recovered hash value as the key, as follows:

$$D_{h(F^i)}(E_{h(F^i)}(F^i)) = F^i.$$

The file block F^i is now in an unencrypted format and ready to be read by the authorized user (i.e., operation 1010).

Writing A File

Fig. 11 shows a process 1100 for writing to or modifying one of the file blocks of a large file without affecting other blocks of the data stream. Once again, it is assumed that the file is of sufficient size to have a tree structure 408 stored in the metadata stream 404 and that the target block is a block F^i . The process 1100 can be implemented in software and will be described with reference to components in the computing device 200 in Fig. 2 and the file format in Figs. 4-6.

At operation 1102, the computing device modifies a portion of the file contained in block F^i , to create a file block $F^{i'}$. Modifying the data renders the previously computed hash value inaccurate and hence unusable. Accordingly, at operation 1104, the control module 220 calls the hash module 226 to compute a new hash value of the modified block, or $h(F^{i'})$. At operation 1106, the control module 220 calls the cryptographic engine 224 to encrypt the modified file block $F^{i'}$ using the new hash value, or $E_{h(F^{i'})}(F^{i'})$. The new encrypted block replaces the pre-modified encrypted block in the data stream 402.

These changes to the file block also affect a portion of the tree structure 408 stored in the metadata stream 404. At operation 1108, the tree builder 230 recreates a new leaf node $L^{i'}$ that is associated with the modified block. The tree builder also recreates any intermediate nodes that reference the new leaf node

1 (either directly or indirectly) as well as the root node (i.e., operation 1110). At
2 operation 1112, the tree builder optionally re-signs the header 406, per-user
3 information 410 and root using the last writer's signature, or using the lock-secret
4 key technique described above

6 **Signed Manifest of File Modifications**

7 In the signed form of the file format, a digital signature is applied to the
8 header 406, per-user information 410 and root node after every modification to the
9 file. This is illustrated, for example, as operation 1112 in the file write process
10 1100 of Fig. 11. The advantage of the unsigned file format over the signed file
11 format is that the writer of a file does not need to compute a digital signature when
12 closing the file after writing to it. Since digital signatures are computationally
13 expensive, this can be a significant savings if file writes are performed frequently.
14 When using the unsigned file format, instead of signing the file, the writer merely
15 sends the file's hash value to the directory servers that implement the directory in
16 which the file is stored. When another machine wishes to verify the contents of a
17 file, it cannot check a signature in the file, since there is no signature in the file to
18 check. The verifying machine thus needs to obtain verification information from
19 some source that is external to the file.

20 One such source is the directory servers that implement the directory in
21 which the file is stored. Since the directory servers store a copy of the file's hash
22 value, they can provide this value to the verifying machine, and the verifying
23 machine can compare this stored hash value to the computed hash value of the file.
24 The disadvantage of this approach is that it requires contacting and trusting the
25 directory servers. The trust issue is not particularly important, since the directory

1 servers already have to be trusted with version information and writer
2 authorizations. However, contacting the directory servers for every file
3 verification can place a significant additional load on these machines, so it is
4 beneficial to avoid this contact if possible.

5 Therefore, the present invention includes a mechanism by which the writer
6 of a file can provide file authentication information to a verifying machine without
7 having to compute a new digital signature every time a written file is closed.
8 Periodically, the writer compiles a list of the hash values of all files that have been
9 written over a recent interval, computes a hash of the list, and signs the hash. This
10 signed list of hash values is known as a manifest, analogous to a shipping manifest
11 that enumerates the items in a shipment. The advantage of using a signed manifest
12 is that the writer need only perform a single signature computation in order to
13 authenticate the writes to multiple files, rather than having to compute a separate
14 signature for each file, as it would for the signed file format.

15 The writing machine can then send the signed manifest, along with one or
16 more of the files that have been written, to a machine that wants a copy of the
17 files. The receiving machine can verify that the signature of the hash of the
18 manifest is valid, that the hash of manifest is valid, and that the file hash in the
19 manifest corresponds to the hash of the file that it is interested in. The verifying
20 machine needs to know the list of authorized writers to the file, which it must
21 obtain from the directory servers, but this list is generally not modified as
22 frequently as the contents of the file, so the load on the directory servers from
23 propagating updates to the authorized writer list is significantly lower than the
24 load from providing a hash value for every new version of a file.

Fig. 12 shows a process 1200 for producing a signed manifest of modifications, and Fig. 13 illustrates an exemplary signed manifest. The process 1200 can be implemented in software and will be described with reference to components in the computing device 200 in Fig. 2 and the exemplary signed manifest in Fig. 13.

At operation 1202, the computing device modifies one or more files. This step is typically performed separately for each file, and depending upon the file size and the scope of the modifications, the control module 220 may invoke one or more of the segmenter 222, the crypto engine 224, the hash module 226, and the tree builder 230 in order to update the file metadata. At operation 1204, the control module 220 calls the hash module 226 to compute a new hash value of each modified file. This step is typically performed separately for each file and in conjunction with the writing of the new data to the file. The control module 220 collects the hash values of every modified file in a manifest (i.e., operation 1206).

Fig. 13 shows an exemplary manifest 1300. It includes a collection of entries 1302-1306 of modified files. Each entry contains both a file number (i.e., the file number 412 in the file header 406) and the hash of file. The file number specifies to which file a particular hash applies. Also, the manifest 1300 includes a magic number header 1308 at the beginning that helps ascertain what is being signed. This is a different magic number than the one found at the beginning of the file header.

After collecting a set of changes, the control module 220 invokes the hash module 226 to compute a hash of the contents of the manifest (i.e., operation 1208), and then it invokes the signing/verifying module 228 to sign the hash of the manifest using the last writer's private key (i.e., operation 1210). This is

1 represented as the signature 1310 in Fig. 13, which covers the entire manifest. By
2 signing the manifest, the file system can verify the user who modified the files in
3 the manner outlined in the manifest. The timing as to when a manifest is created
4 varies depending upon implementation requirements.

5 The manifest may be subsequently verified by initially verifying the
6 signature 1310. If the signature is valid, the file hash contained in the manifest is
7 compared to the hash of the file. If the two match, the verifier will then examine
8 the revision number in the file. Action is only taken if the revision number in the
9 file is bigger than the biggest revision number that the verifier has ever seen for
10 that file. With this last evaluation, the verifier prevents malicious/malfunctioning
11 machines from pushing stale versions of files to replica sites.

12 13 **Sparse Files**

14 A sparse file is a file whose logical size is greater than its physical size,
15 potentially possessing large ranges with no data whatsoever. Such ranges are said
16 to be “unallocated” as distinguished from “allocated” ranges that contain actual
17 data. Logically, unallocated regions of a file contain zero data, but there is no
18 physical storage associated with these regions. Sparse files are used in many
19 different environments (e.g., database logging) and are known in the art. The file
20 format described above can be used to support sparse files; however, it may
21 consume an inefficiently large amount of storage space. This section describes
22 modifications to the file format that greatly reduce this inefficiency. No fields are
23 added, removed, or rearranged. However, for efficient handling of sparse files, a
24 few changes can be made to the way values in certain fields are calculated.

1 If one were to store a sparse file using the file format 400 illustrated in Fig.
2 4, the allocated size of the metadata stream 404 will be proportional to the total
3 size of the primary data stream 402 that includes vast ranges with no content (i.e.,
4 the logical file size), rather than proportional to the allocated size of the primary
5 data stream that includes only the content portions of the sparse file (i.e., the
6 physical file size). One significant case of sparse file usage is for circular logging
7 using truncate-from-tail. With the file format 400, the metadata 404 would grow
8 linearly as the log is written, but it could not be truncated as the log is truncated.
9 Therefore, a file that is expected to be constant in size would actually grow
10 linearly without bound.

11 The modifications described below adapt the file format so that (1)
12 unallocated primary-stream plaintext is represented by unallocated primary-stream
13 ciphertext, and (2) unallocated primary-stream data produces corresponding
14 metadata of all zeroes, thereby enabling use of sparse file allocation for the
15 metadata stream. The modifications may not completely eliminate the
16 inefficiencies in allocation of the metadata stream, but the allocated metadata size
17 will always be proportional to the allocated primary stream size and at worst
18 logarithmically related to the total primary stream size. Further, in the circular
19 logging case, the modifications allow the metadata to be truncated as the primary
20 data stream is truncated.

21 Generally, the modifications differentiate the unallocated regions of a
22 sparse file that contain no real content from the allocated portions of the sparse
23 file. Once differentiated, the file system can deallocate the non-content portions.
24 In one implementation, the file system creates a new one-way hash function $g(x)$,
25 as follows:

if $x = 0$

$g(x) = 0$

else

$g(x) = h(x)$

where $h(x)$ is the standard one-way hash function specified in the file format above. The hash function $g(x)$ has the property that data of all zeroes hashes to a hash value of all zeroes.

One other modification is made to the leaf nodes of the tree structure 408 that are associated with file blocks in unallocated ranges that contain no ciphertext. Each leaf node associated with such file blocks is modified such that the access value 604 (i.e., the encrypted hash (irrespective of the encryption key) of the nonexistent plaintext) is set to zero and the verification value 606 (i.e., the hash of the nonexistent ciphertext) is set to zero. Following this adjustment to the leaf nodes, the tree structure 408 is constructed using the hash function $g(x)$ so higher-order intermediate nodes in the tree for zero-value leaf nodes will also be zero: $g(0) = 0$.

In this manner, the file system need only allocate ranges for metadata blocks that contain non-zero data, which will be those that correspond to allocated primary stream data. Files that have large ranges of unallocated blocks, such as circular-logging files, will have correspondingly large portions of zero-value metadata. The file system can then simply deallocate this metadata without changing its semantics.

1 The modifications discussed in this section do not compromise file-write
2 security. It may seem to, since write security rests on the non-invertibility of the
3 one-way hash function, and the non-invertibility in the special case of zero-value
4 data has been compromised. However, all that has been forfeited is that an
5 attacker can trivially compute the hash of zero-value data, but an attacker could
6 have easily computed this value anyway simply by performing the hash
7 computation.

8 Note that the stored hash value for unallocated ciphertext is zero, whereas
9 the stored hash value for zero-value ciphertext is $h(0)$, and the stored hash value
10 for ciphertext corresponding to zero-value plaintext is $h(E(0))$. Thus, the tree of
11 hashes distinguishes between all-zero primary-stream blocks and unallocated
12 primary-stream blocks. This prevents an attacker from substituting one of these
13 for the other without detection by the directory servers or storage servers. Such a
14 substitution has the ability to affect application behavior, since applications can
15 query the set of allocated ranges in a file.

16 One potential downside is that the modifications for supporting sparse files
17 do slightly compromise file-read security because it allows an attacker with no
18 access to cryptographic keys to determine ranges in a file that are unallocated.
19 However, this is not believed to result in a significant information leak.

21 Conclusion

22 Although the invention has been described in language specific to structural
23 features and/or methodological steps, it is to be understood that the invention
24 defined in the appended claims is not necessarily limited to the specific features or
25

steps described. Rather, the specific features and steps are disclosed as preferred forms of implementing the claimed invention.